

Approximate Bayesian Computation via Regression Density Estimation

Y. FAN,¹ D. J. NOTT² and S. A. SISSON¹

ABSTRACT

Approximate Bayesian computation (ABC) methods, which are applicable when the likelihood is difficult or impossible to calculate, are an active topic of current research. Most current ABC algorithms directly approximate the posterior distribution, but an alternative, less common strategy is to approximate the likelihood function. This has several advantages. First, in some problems, it is easier to approximate the likelihood than to approximate the posterior. Second, an approximation to the likelihood allows reference analyses to be constructed based solely on the likelihood. Third, it is straightforward to perform sensitivity analyses for several different choices of prior once an approximation to the likelihood is constructed, which needs to be done only once. The contribution of the present paper is to consider regression density estimation techniques to approximate the likelihood in the ABC setting. Our likelihood approximations build on recently developed marginal adaptation density estimators by extending them for conditional density estimation. Our approach facilitates reference Bayesian inference, as well as frequentist inference. The method is demonstrated via a challenging problem of inference for stereological extremes, where we perform both frequentist and Bayesian inference.

Keywords: Approximate Bayesian computation; Copulas; Likelihood-free inference; Multivariate density estimation; Regression density estimation.

1 Introduction

Approximate Bayesian computation (ABC) methods (commonly described as “likelihood-free” methods) have been attracting increasing research interest as a viable procedure for performing Bayesian inference in the presence of computationally intractable likelihood functions. Initially popular in the biological sciences (Beaumont et al. 2002; Luciani et al. 2009; Ratmann et al. 2007; Ratmann et al. 2009), ABC has now found application in a wide range of areas – see the reviews by Beaumont (2010), Csilléry et al. (2010) and Sisson and Fan (2011).

ABC methods require the ability to simulate data from the intractable likelihood, $L(\cdot|\theta)$, for a given parameter vector θ . An integral part of all existing algorithms involves sampling θ from some distribution (typically the prior, $\pi(\theta)$) and comparing simulated data $\mathbf{x} \sim L(\cdot|\theta)$ with the observed data \mathbf{x}_0 . Posterior samples approximately from

¹School of Mathematics and Statistics, University of New South Wales, Sydney, 2052, AUSTRALIA

²Dept. of Statistics and Applied Probability, National University of Singapore, SINGAPORE, 117546

$\pi(\boldsymbol{\theta}|\mathbf{x}_0)$ are then obtained by weighting $\boldsymbol{\theta}$ according to some weighting function or kernel $K_\delta(\|\mathbf{x} - \mathbf{x}_0\|)$ with scale parameter $\delta \geq 0$. Accordingly, samples $\boldsymbol{\theta}$ for which $\mathbf{x} \approx \mathbf{x}_0$ receive larger weights than those where \mathbf{x} is very different to \mathbf{x}_0 .

More formally, ABC algorithms produce samples from an augmented posterior distribution

$$\pi_{ABC}(\boldsymbol{\theta}, \mathbf{x}|\mathbf{x}_0) \propto K_\delta(\|\mathbf{x} - \mathbf{x}_0\|)L(\mathbf{x}|\boldsymbol{\theta})\pi(\boldsymbol{\theta}). \quad (1)$$

It is easy to see that as $\delta \rightarrow 0$, only $\mathbf{x} \approx \mathbf{x}_0$ is given non-trivial weight by K_δ , and the ABC posterior converges to the true posterior distribution, $\pi(\boldsymbol{\theta}|\mathbf{x}_0)$. However, $\delta = 0$ practically equates to repeated simulation until $\mathbf{x} = \mathbf{x}_0$ is matched exactly for each $\boldsymbol{\theta}$, which is computationally impractical in general. As such, for some $\delta > 0$, a necessarily approximate posterior is obtained. Regression-adjustment strategies have been proposed to post-process the approximate posterior samples to correct the error in the ABC posterior (e.g. Beaumont et al. 2002; Blum and Francois 2010; Nott et al. 2012). Dimension reduction techniques are also typically used to reduce the overall computational overhead. Here, the vectors, \mathbf{x} , are replaced by lower-dimensional sufficient or summary statistics, $\mathbf{S} \equiv \mathbf{S}(\mathbf{x})$. Readers should refer to Blum et al. (2012), Fearnhead and Prangle (2012) and Robert et al. (2011) for further discussion of the use and elicitation of summary statistics in ABC.

In this paper, we present an alternative ABC approach based on constructing a direct approximation to the likelihood. The approximation is based on samples $\boldsymbol{\theta}^{(i)}$ (not necessarily from the prior), and the corresponding summaries $\mathbf{S}^{(i)}$, simulated from the intractable likelihood. Such an approach is attractive in several ways. Firstly, it is sometimes easier to approximate the likelihood than to approximate the posterior directly. Secondly, separate estimation of the likelihood is useful for performing purely likelihood based analyses, such as maximum likelihood estimation. This may be more attractive to researchers more familiar with frequentist inference. Thirdly, direct approximation of the likelihood is useful for diagnosing prior-likelihood conflict, or for running many analyses with different priors in a sensitivity analysis, as the likelihood approximation only needs to be constructed once. A further example of a Bayesian use for a separate estimate of the likelihood is the construction of credible regions based on contours of the likelihood – these have a robust Bayes interpretation in terms of posterior probability content being minimally sensitive to perturbations of the prior (Wasserman 1989). The method we propose here provides an explicit analytical expression for the likelihood function. It also does not require simulation from the prior, which is typically required by many ABC algorithms (see e.g. Pritchard et al. 1999; Tavaré et al. 1997; Sisson et al. 2007). This is useful where there is interest in using reference priors.

An early discussion of the concept of directly approximating an intractable likelihood function using simulated data sets was given by Diggle and Gratton (1984), predating the current wave of activity in ABC research. More recently, several researchers have considered direct likelihood approximations within ABC. Leuenberger and Wegmann (2010) developed an approach that involves embedding the regression adjustment of Beaumont et al. (2002) into a likelihood approximation based on a generalised linear model. Wood (2010) considers a “synthetic likelihood” which involves estimating a normal distribution for the summary statistics, \mathbf{S} , with mean and covariance depending on $\boldsymbol{\theta}$. The approximation of the mean and covariance is performed within a Markov chain Monte Carlo (MCMC) scheme, and through clever choice of summary statistics and quantile transformations it may be possible to improve the approximation to normality. The approach we

propose is also related to that introduced by Bonassi et al. (2011), although this method does not directly approximate the likelihood. They propose to use mixtures of multivariate normals to estimate a joint distribution for $(\boldsymbol{\theta}, \mathbf{S})$ where the summary statistics are simulated from an informative prior, and then condition on the observed summary statistic in the estimated mixture model. Appropriate localization and choice of summary statistics can help to improve the efficiency of the approach.

We note that a direct approximation of the likelihood function is available by Monte Carlo integration (e.g. Sisson et al. 2008; Sisson and Fan 2011). If we consider marginalising the distribution of the augmented posterior (1) with respect to the auxiliary data \mathbf{x} , we then have the marginal posterior

$$\pi_{ABC}(\boldsymbol{\theta}|\mathbf{x}_0) \propto \pi(\boldsymbol{\theta}) \int K_\delta(\|\mathbf{x} - \mathbf{x}_0\|) L(\mathbf{x}|\boldsymbol{\theta}) d\mathbf{x} \approx \frac{\pi(\boldsymbol{\theta})}{n} \sum_{i=1}^n K_\delta(\|\mathbf{x}^{(i)} - \mathbf{x}_0\|) \quad (2)$$

where $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}$, are samples from $L(\mathbf{x}|\boldsymbol{\theta})$. The summation term is a simple Monte Carlo estimate of the likelihood of \mathbf{x}_0 at the point $\boldsymbol{\theta}$. Clearly, this estimate is a function of δ , and so posterior inferences will still require regression-adjustment post-processing. It is also a pointwise estimate, and so it must be effectively re-estimated for each likelihood evaluation in an analysis. Our approach results in a stand-alone functional expression for the likelihood function at $\delta = 0$. However, in the context of this paper, the estimator (2) could usefully serve as a goodness-of-fit diagnostic.

Our proposed approach to conditional density estimation for ABC (that is, estimating the distribution of $\boldsymbol{\theta}|\mathbf{S}$ given the joint samples $(\boldsymbol{\theta}, \mathbf{S})$) builds on recent work on flexible multivariate density estimation – in particular the marginal adaptation method of Giordani et al. (2012). The basis of their approach argues that estimation of univariate marginal distributions is easier than estimating a full multivariate density estimation directly. Hence, a fruitful strategy is to adapt a multivariate density estimate to have given marginals, which are more precisely estimated individually. In this paper, we develop a strategy for implementing this approach for the related problem of conditional density estimation, in the context of likelihood estimation in ABC.

In Section 2, we develop the estimator for the intractable likelihood function, and Section 3 outlines some connections with existing ABC approaches. Section 4 discusses a challenging example of inference for stereological extremes and Section 5 concludes.

2 ABC via multivariate regression density estimation

Our goal is to obtain an estimate of the likelihood $L(\mathbf{S}|\boldsymbol{\theta})$, denoted by $\hat{L}(\mathbf{S}|\boldsymbol{\theta})$, and then approximate the intractable posterior distribution by

$$\hat{\pi}(\boldsymbol{\theta}|\mathbf{S}_0) \propto \hat{L}(\mathbf{S}_0|\boldsymbol{\theta})\pi(\boldsymbol{\theta}),$$

where $\mathbf{S}_0 = \mathbf{S}(\mathbf{x}_0)$ denotes the summary statistics from the observed data. Once $\hat{L}(\mathbf{S}_0|\boldsymbol{\theta})$ is obtained, drawing samples from $\hat{\pi}(\boldsymbol{\theta}|\mathbf{S}_0)$ can be achieved by any standard posterior simulation algorithm e.g. Markov chain Monte Carlo. We assume that the vector of summary statistics, $\mathbf{S} = (S^1, \dots, S^k)^\top$, are given, with $\dim(\mathbf{S}) = k$ and that $\dim(\boldsymbol{\theta}) = d$.

The likelihood estimate is constructed from samples $(\mathbf{S}^{(i)}, \boldsymbol{\theta}^{(i)})$ for $i = 1, \dots, N$, where $\boldsymbol{\theta}^{(i)} \sim h(\boldsymbol{\theta})$ and $\mathbf{S}^{(i)} = \mathbf{S}(\mathbf{x}^{(i)})$ with $\mathbf{x}^{(i)} \sim L(\cdot|\boldsymbol{\theta}^{(i)})$. The distribution function $h(\boldsymbol{\theta})$ determines the region of parameter space in which we require $\hat{L}(\mathbf{S}_0|\boldsymbol{\theta})$ to be a good approximation of $L(\mathbf{S}_0|\boldsymbol{\theta})$. This must necessarily include the region of high posterior density. One option, where this is convenient, is to perform an initial pilot ABC analysis with a large value of δ to broadly identify the high posterior density region, and then define $h(\boldsymbol{\theta})$ to be proportional to the prior $\pi(\boldsymbol{\theta})$ in this region (Fearnhead and Prangle 2012).

We first consider flexible estimation of the individual marginal distributions $S^j|\boldsymbol{\theta}$ for each $j = 1, \dots, k$. Estimation of marginals is typically easier than estimation of the joint distribution due to the lower dimensionality, and there are many different methods available for this purpose. Here we adopt a mixture of experts approach (Jacobs et al. 1991; Jordan and Jacobs 1994). Mixtures of experts models are mixtures of regression models where both parameters in the component response distributions and mixing weights are allowed to vary with covariates. In our case, for each S^j , we express the marginal distribution of $S^j|\boldsymbol{\theta}$ as a mixture of J normal distributions

$$f_j(S^j|\boldsymbol{\theta}) = \sum_{\ell=1}^J w_{j\ell}(\boldsymbol{\theta}) N(\mu_{j\ell}(\boldsymbol{\theta}), \sigma_{j\ell}^2(\boldsymbol{\theta})) \quad (3)$$

where

$$w_{j\ell}(\boldsymbol{\theta}) = \frac{\exp(\xi_0^{j\ell} + \xi^{j\ell\top} \boldsymbol{\theta})}{\sum_{k=1}^J \exp(\xi_0^{jk} + \xi^{jk\top} \boldsymbol{\theta})}.$$

Here $\xi_0^{j\ell} \in \mathbb{R}$ and $\xi^{j\ell} \in \mathbb{R}^d$, for $j = 1, \dots, k$, $\ell = 1, \dots, J$ with $\xi_0^{j1} = 0$, $\xi^{j1} = \mathbf{0}$ for identifiability, and

$$\mu_{j\ell}(\boldsymbol{\theta}) = \beta_0^{j\ell} + \boldsymbol{\beta}^{j\ell\top} \boldsymbol{\theta}, \quad \log(\sigma_{j\ell}^2(\boldsymbol{\theta})) = \gamma_0^{j\ell} + \boldsymbol{\gamma}^{j\ell\top} \boldsymbol{\theta},$$

with $\beta_0^{j\ell}, \gamma_0^{j\ell} \in \mathbb{R}$ and $\boldsymbol{\beta}^{j\ell}, \boldsymbol{\gamma}^{j\ell} \in \mathbb{R}^d$. For a sufficiently large sample size (N) and number of components (J), the model (3) can approximate the marginal distribution of $S^j|\boldsymbol{\theta}$ arbitrarily well. Methods for fast and parsimonious fitting of (3) can be found in Nott et al. (2012) and Tran et al. (2012).

For the estimation of the joint dependencies, we propose a mixture of Gaussian copulas approach similar to that proposed by Giordani et al. (2012) for density estimation. More precisely, for joint samples of $(\mathbf{S}^{(i)}, \boldsymbol{\theta}^{(i)})$ we firstly transform the fitted margins of S^j to be standard normal, so that

$$U^j = \Phi^{-1}(F_j(S^j|\boldsymbol{\theta})), \quad (4)$$

where $\mathbf{U} = (U^1, \dots, U^k)$, F_j denotes the cumulative distribution function for the fitted density $f_j(\cdot)$ in (3) and Φ denotes the standard normal cumulative distribution function. We then fit a mixture of multivariate normals to the joint samples $(\mathbf{U}, \boldsymbol{\theta})$, and obtain the conditional distribution of $\mathbf{U}|\boldsymbol{\theta}$ in this mixture (see e.g. Norets and Pelenis 2012). In more detail, the joint density of $(\mathbf{U}, \boldsymbol{\theta})$ is modelled via a mixture of multivariate normal distributions

$$g(\mathbf{U}, \boldsymbol{\theta}) = \sum_{\ell=1}^L \omega_{\ell} N_{k+d}(\boldsymbol{\nu}_{\ell}, \boldsymbol{\Sigma}_{\ell}), \quad (5)$$

where $\sum_{\ell=1}^L \omega_\ell = 1$ and $\boldsymbol{\nu}_\ell, \boldsymbol{\Sigma}_\ell$ denote the mean and covariance matrices of the normal mixture components. The conditional density $g(\mathbf{U}|\boldsymbol{\theta})$ is then given by

$$g(\mathbf{U}|\boldsymbol{\theta}) = \sum_{\ell=1}^L \omega_\ell^c N_k(\boldsymbol{\nu}_\ell^c, \boldsymbol{\Sigma}_\ell^c), \quad (6)$$

where $\boldsymbol{\nu}_\ell^c$ and $\boldsymbol{\Sigma}_\ell^c$ are the mean and covariance matrix of the conditional distribution given $\boldsymbol{\theta}$ in the ℓ -th mixture component, $N_{k+d}(\boldsymbol{\nu}_\ell, \boldsymbol{\Sigma}_\ell)$. The mixing weights for the conditional distribution, ω_ℓ^c , are

$$\omega_\ell^c = \frac{\omega_\ell \phi(\boldsymbol{\theta}, \boldsymbol{\nu}_\ell(\boldsymbol{\theta}), \boldsymbol{\Sigma}_\ell(\boldsymbol{\theta}))}{\sum_{i=1}^L \omega_i \phi(\boldsymbol{\theta}, \boldsymbol{\nu}_i(\boldsymbol{\theta}), \boldsymbol{\Sigma}_i(\boldsymbol{\theta}))}$$

where $\phi(\boldsymbol{\theta}, \boldsymbol{\nu}_\ell(\boldsymbol{\theta}), \boldsymbol{\Sigma}_\ell(\boldsymbol{\theta}))$ denotes the marginal density for $\boldsymbol{\theta}$ in the normal component $N_{k+d}(\boldsymbol{\nu}_\ell, \boldsymbol{\Sigma}_\ell)$.

Our density estimator for the likelihood is then given as

$$\hat{L}(\mathbf{S}|\boldsymbol{\theta}) = g(\mathbf{U}|\boldsymbol{\theta}) \prod_{j=1}^k \frac{f_j(S^j|\boldsymbol{\theta})}{\phi(U^j, 0, 1)}, \quad (7)$$

where $\phi(U^j, 0, 1)$ denotes a standard normal density evaluated at U^j . Note that $\hat{L}(\mathbf{S}|\boldsymbol{\theta})$ does not have exactly $f_j(S^j|\boldsymbol{\theta})$ as its marginal distribution, since the marginal distribution for U^j under $g(\mathbf{U}|\boldsymbol{\theta})$ is not exactly standard normal, unless $L = 1$. For this reason, Giordani et al. (2012) suggest recomputing the distribution of U^j using another normal mixture, which would lead to replacing $\phi(U^j, 0, 1)$ with this distribution in (7). This will improve the quality of the likelihood approximation, albeit at additional computational cost. However in our experience we have found that an acceptable approximation is obtained even without enforcing the marginals $f_j(S^j|\boldsymbol{\theta})$ exactly.

A key component of the above approach is the transformation $(\mathbf{S}, \boldsymbol{\theta}) \rightarrow (\mathbf{U}, \boldsymbol{\theta})$. If the summary statistics are too highly correlated, or if the dependence between the statistics and the parameters is too complex, it will be difficult to capture the joint distribution of $(\mathbf{S}, \boldsymbol{\theta})$ well in a simple mixture model. The primary role of the flexible mixture of experts approach (3) is to construct the transformation $(\mathbf{S}, \boldsymbol{\theta}) \rightarrow (\mathbf{U}, \boldsymbol{\theta})$ to simplify this complexity. However, such modelling can be greatly aided by careful initial selection of summary statistics to achieve near orthogonality and low dependence (e.g. Blum et al. 2012, Fearnhead and Prangle 2012).

A step by step description of how to construct $\hat{L}(\mathbf{S}|\boldsymbol{\theta})$ is:

Step 1 Sample $(\mathbf{S}^{(i)}, \boldsymbol{\theta}^{(i)}) \sim L(\mathbf{S}|\boldsymbol{\theta})h(\boldsymbol{\theta})$ for $i = 1, \dots, N$. Define $\mathbf{S}^{(i)} = (S^{(i)1}, \dots, S^{(i)k})$.

Step 2 For each S^j , fit the mixture model (3) to the marginal samples $(S^{(i)j}, \boldsymbol{\theta}^{(i)})$.

Step 3 Compute the transformations $U^{(i)j} = \Phi^{-1}(F_j(S^{(i)j}|\boldsymbol{\theta}))$, for $j = 1, \dots, k$, $i = 1, \dots, N$. Define $\mathbf{U}^{(i)} = (U^{(i)1}, \dots, U^{(i)k})$.

Step 4 Fit the mixture model (5) to the samples $(\mathbf{U}^{(1)}, \boldsymbol{\theta}^{(1)}), \dots, (\mathbf{U}^{(N)}, \boldsymbol{\theta}^{(N)})$ and derive the model $g(\mathbf{U}|\boldsymbol{\theta})$ from (6).

Step 5 Combine the results in Steps 2, 3 and 4 to obtain the final estimate for the likelihood $\hat{L}(\mathbf{S}|\boldsymbol{\theta})$ via (7).

The estimate, $\hat{L}(\mathbf{S}|\boldsymbol{\theta})$, can then be used as part of any standard Bayesian or frequentist analysis.

3 Connections with other methods

Bonassi et al. (2011) provide a method to directly estimate the posterior density. A mixture of multivariate normals is used to approximate the joint distribution of $(\mathbf{S}, \boldsymbol{\theta})$ based on samples $(\mathbf{S}^{(i)}, \boldsymbol{\theta}^{(i)}) \sim L(\mathbf{S}|\boldsymbol{\theta})\pi(\boldsymbol{\theta})$, where the prior must be informative. Conditioning this mixture on $\mathbf{S} = \mathbf{S}_0$ then provides an estimate of the posterior. To the best of our knowledge, this is the only approach that directly uses non-parametric conditional density estimation techniques in ABC. Note that alternatively conditioning the mixture approximation of $(\mathbf{S}, \boldsymbol{\theta})$ on $\boldsymbol{\theta}$ would result in a mixture of normals conditional density estimate for the likelihood. Similarly, we may derive a direct estimate of the posterior similar to (7) by conditioning on \mathbf{S} rather than $\boldsymbol{\theta}$, although the sampling distribution $h(\boldsymbol{\theta})$ must then equal the prior $\pi(\boldsymbol{\theta})$. Arguments for estimating likelihoods, rather than posterior distributions are provided in the Introduction.

Giordani et al. (2012) discuss the performance of the mixture of multivariate normals density estimator, and observe that it can be difficult to estimate the marginals well using this approach. Our conditional estimator (7) uses a mixture of normals to estimate the joint distribution of transformed data. The benefit of the transformation $(\mathbf{S}, \boldsymbol{\theta}) \rightarrow (\mathbf{U}, \boldsymbol{\theta})$ is that it is simpler to estimate the joint distribution of $(\mathbf{U}, \boldsymbol{\theta})$ than that of $(\mathbf{S}, \boldsymbol{\theta})$. The mixture of normals model also provides greater flexibility over Gaussian copula models which would simply fit a normal distribution to the transformed data.

There are clear benefits in using more flexible density estimation methods. For example, the flexibility in the mixture of experts model (3) to describe the relationship between S^j and $\boldsymbol{\theta}$, means that good transformations $(\mathbf{S}, \boldsymbol{\theta}) \rightarrow (\mathbf{U}, \boldsymbol{\theta})$ can be constructed for a wide range of $\boldsymbol{\theta}$ vectors drawn from $h(\boldsymbol{\theta})$. As such, we can obtain good approximations to $L(\mathbf{S}|\boldsymbol{\theta})$ for a wide range of $\boldsymbol{\theta}$ as long as the regression models in (3) are adequate and enough samples $(\mathbf{S}^{(i)}, \boldsymbol{\theta}^{(i)})$ are available. Clearly, well-chosen summary statistics can also make the modelling easier, especially when they behave nearly linearly with the parameters. See Blum et al. (2012) for a review of dimension reduction methods for ABC.

Finally, we note that if the initial sample $(\mathbf{S}^{(i)}, \boldsymbol{\theta}^{(i)})$ is obtained via a pilot ABC analysis with a moderately large kernel scale parameter δ , then our approach can be viewed as another form of ABC posterior adjustment method. This is in the sense that the initial samples are used to estimate the likelihood at the point $\delta = 0$, from which (new) samples are subsequently obtained. See e.g. Leuenberger and Wegmann (2010) and other regression adjustment postprocessing methods (Beaumont et al. 2002; Blum and Francois 2010; Nott et al. 2012) for further information.

4 Example: Inference for stereological extremes

We now present a reanalysis of a stereological problem studied by Bortot et al. (2007) focusing on the production of clean steels. Here inference is required on the size and

intensity of the largest microscopic particle inclusions in a 3-dimensional block of steel, based on inclusion cross sections observed in a 2-dimensional planar slice. The observed data consist of 112 inclusions intersecting the planar slice, with cross sectional diameters above a measurement threshold of $\nu_0 = 5\mu m$. We follow Bortot et al. (2007) and consider two models. The first assumes a Poisson process with rate λ for inclusion centres, and a size distribution for the inclusions, which are assumed to be spherical, based on univariate extreme value theory. Here, conditional on exceeding the threshold, ν_0 , inclusion sizes are assumed to follow a generalised Pareto distribution with scale and shape parameters $\sigma > 0$ and ξ . In this setting it is possible to perform standard Bayesian inference, and so the correct posterior distribution is available. The second model is the same as the first, but assumes that the inclusions are ellipsoidal. For each model we have $\boldsymbol{\theta} = (\lambda, \sigma, \xi)$.

The analysis of Bortot et al. (2007) took the number of inclusions, and 112 equally spaced quantiles, $q_{(1)}, \dots, q_{(112)}$, as summary statistics. Here, as the quantiles will be highly correlated, we instead use $S^j = \log(q_{(j+1)} - q_{(j)})$ for $j = 1, \dots, 111$, where S^{112} is the log number of inclusions. In order to determine the distribution $h(\boldsymbol{\theta})$ in regions of non-negligible posterior density, we found it convenient to first draw samples from $I(\|\mathbf{S} - \mathbf{S}_0\| \leq \delta)L(\mathbf{S}|\boldsymbol{\theta})U(\boldsymbol{\theta})$, where $\|\cdot\|$ denotes squared Euclidean distance, for a large value of $\delta = 20$, where $U(\boldsymbol{\theta})$ is uniform over a moderately large region of parameter space expected to contain the posterior. We then defined

$$h(\boldsymbol{\theta}) = N(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}})I\left((\boldsymbol{\theta} - \hat{\boldsymbol{\mu}})^\top \hat{\boldsymbol{\Sigma}}^{-1}(\boldsymbol{\theta} - \hat{\boldsymbol{\mu}}) \leq 3d\right)$$

as a truncated normal distribution, where $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{\Sigma}}$ denote the mean and covariance of the $\boldsymbol{\theta}$ -component of these samples. In this manner, we obtained a final sample $(\mathbf{S}^{(i)}, \boldsymbol{\theta}^{(i)}) \sim L(\mathbf{S}|\boldsymbol{\theta})h(\boldsymbol{\theta})$ of size 5,000 for each model.

In the following, we demonstrate the performance of the regression density estimation methodology using the above summary statistics (corresponding to conditional density estimation in 115 dimensions), as well as the “semi-automatic” summary statistics of Fearnhead and Prangle (2012), which provide one summary statistic per parameter (i.e. 6-dimensional density estimation). In addition to standard ABC posterior inference we also perform a prior sensitivity analysis and a maximum-likelihood based inference.

4.1 Results

To fit the mixture of experts marginal distributions f_j (3), we adopted the variational Bayes approach of Nott et al. (2012) and used the variational lower bound to select the number of mixture components. For both spherical and ellipsoidal models we found $J = 3$ mixture components worked well for each of the 112 summaries, although $J = 6$ components were required for the 3 semi-automatic statistics. Note that the fitting of the marginal models, f_j , is highly parallelizable, and is accordingly suitable for applications in high dimensions. Figure 1 shows the relationship between the parameters (λ, σ, ξ) and the three semi-automatic summary statistics, as observed in the sample $(\mathbf{S}^{(i)}, \boldsymbol{\theta}^{(i)}) \sim L(\mathbf{S}|\boldsymbol{\theta})h(\boldsymbol{\theta})$ for the ellipsoidal model. There are visibly strong relationships both within the summary statistic and parameter vectors, and between them. Figure 2 displays the same information following the transformation $(\mathbf{S}, \boldsymbol{\theta}) \rightarrow (\mathbf{U}, \boldsymbol{\theta})$. Clearly, the transformation has greatly simplified the underlying dependence structure. As a result, while the joint distribution of $(\mathbf{S}, \boldsymbol{\theta})$ could be reasonably well modelled by a mixture of multivariate normal distributions (e.g. Bonassi et al. 2011), a better fit is likely to be ob-

tained by alternatively fitting to $(\mathbf{U}, \boldsymbol{\theta})$.

Figure 3 shows the fitted marginal density estimate, $f_j(S^j|\boldsymbol{\theta})$, for $S^j, j = 5, 40, 80, 110$ under the spherical model, evaluated at $\boldsymbol{\theta}_0 = (\lambda_0, \sigma_0, \xi_0) = (30, 1.5, -0.05)$. The histogram corresponds to the true density, obtained from data generated under the model $L(\mathbf{S}|\boldsymbol{\theta}_0)$ at the point $\boldsymbol{\theta}_0$, and the solid line denotes the fitted density. The plots indicate a very good fit, and this was typical of all other summary statistics, S^j , and conditioned parameters, $\boldsymbol{\theta}_0$, examined. These results indicate that marginal densities can be very well estimated. Indeed, marginal densities are much easier to estimate in general than joint densities.

We used the R package `Mclust` to fit the joint distribution $g(\mathbf{U}, \boldsymbol{\theta})$ (5) to the transformed samples, using $L = 5$ mixture components for both spherical and ellipsoidal models. Model selection for the number of mixture components is typically performed using the BIC. However care should be taken in high dimensions, as the dimension of $(\mathbf{U}, \boldsymbol{\theta})$ can grow quickly, and the BIC tends to overpenalise the number of parameters. Giordani et al. (2012) use the variational Bayes approximation method to fit the mixture model, and advocate the use of the variational lower bound for model selection.

The histograms in Figure 4 shows the regression density estimates of the marginal posteriors for the spherical inclusion model, based on 500,000 MCMC samples, obtained under the prior specification $\log \lambda \sim N(0, 100^2)$, $\sigma \sim Ga(0.01, 0.0001)$ and $\xi \sim N(0, 100^2)$. The top panels are based on using 112 summary statistics, with the middle panels using the 3 semi-automatic summary statistics. The solid line indicates a density estimate of the true posterior marginal distribution derived from MCMC output using the known likelihood for the spherical inclusions model. The quantile-quantile plots compare the posterior estimated with the semi-automatic statistics with the true marginals. Even when constructing regression density estimates of the likelihood in 115 dimensions, reasonably good marginal parameter estimates are obtained, although they are clearly not perfect. The estimates improve considerably when only using 3 summary statistics, as the regression density estimation is then performed in only 6 dimensions. Clearly, the regression density approach has worked well, even in high dimensions. Recall that these approximations are obtained based on flexible modelling with only 5,000 well placed samples $(\mathbf{S}^{(i)}, \boldsymbol{\theta}^{(i)})$ (though there are some computational overheads to identify where to place these). This compares favourably with earlier analyses of these data (Bortot et al. 2007) with the number of simulations in the tens of millions.

The histograms in Figure 5 display the marginal posterior regression density estimates for the ellipsoidal model. Again, the top panels are based on using 112 summary statistics, with the bottom panels using the 3 semi-automatic summary statistics. The solid lines represent the estimate of the posterior distribution obtained using the MCMC-based ABC method of Bortot et al. (2007) with kernel scale parameter $\delta = 3.3$. Under the spherical model, this method is able to correctly reproduce the true posterior for $\delta = 3.3$ (or lower). Bortot et al. (2007) then assumed that this approach could correctly reproduce the true posterior for the ellipsoidal model, although they were unable to verify their results by reducing δ further. From Figure 5, as with the spherical inclusions model, while the marginal posterior distributions under the regression density approach using the 112 summary statistics appear less accurate than when using the semi-automatic statistics, they still have broadly the right shape and location. Of course, because of the lower-dimensional regression density modelling, when using the

3-dimensional semi-automatic statistics, the resulting density estimates are closer to the MCMC-based estimate of Bortot et al. (2007), albeit obtained with a far smaller computational overhead. In summary, these results appear to broadly support the previous study of the ellipsoidal model by Bortot et al. (2007), although there is some indication that the posterior marginals obtained by Bortot et al. (2007) may be too wide, especially for λ . This would be consistent with the regression density estimate's construction at $\delta = 0$ compared with that of $\delta = 3.3$ for Bortot et al. (2007).

As previously discussed, once the estimate of $L(\mathbf{S}|\boldsymbol{\theta})$ is obtained, it is then a simple matter to perform a range of likelihood-based inferences that are not immediately available with standard ABC methods. Table 1 provides numerical estimates of the mean and the 0.025 and 0.975 quantiles of the marginal posterior for each parameter under various models and analyses, based on the semi-automatic summary statistics. The estimates $(\lambda_{mle}, \sigma_{mle}, \xi_{mle})$ denote the maximum likelihood estimates based on $\hat{L}(\mathbf{S}|\boldsymbol{\theta})$, obtained by maximising the analytical expression for the approximate likelihood using the `optim` function in R. These frequentist estimates are consistent with their Bayesian counterparts in Table 1 due to the uninformative prior implementation. The estimates $(\lambda_{ABC}, \sigma_{ABC}, \xi_{ABC})$ indicate those obtained using the MCMC-based ABC method of Bortot et al. (2007) with kernel scale parameter $\delta = 3.3$. These analyses use the 112-dimensional summary statistics. For the spherical model, this ABC posterior is known to coincide with the true posterior, and this is in evidence in the results, although the upper tail for λ_{ABC} is estimated slightly shorter. As with Figure 5, the marginal posteriors for the Bortot et al. (2007) analysis are wider than those obtained from $\hat{L}(\mathbf{S}|\boldsymbol{\theta})$.

Finally, the estimates $(\lambda_a, \sigma_a, \xi_a)$ correspond to those obtained under the regression density approach for different choices of the prior for $\xi \sim N(0, a^2)$, for $a = 100, 1, 0.5$ and 0.1 . The original analysis of Bortot et al. (2007) with $a = 100$ did not incorporate any information about the upper-tail heaviness of inclusion sizes: $\xi > 2$ is physically unlikely, as this would result in extremely heavy tails; $\xi < -1$ is also unlikely as this would result in a very clear maximum inclusion size truncation. It is computationally trivial to implement this prior sensitivity analysis once the estimate $\hat{L}(\mathbf{S}|\boldsymbol{\theta})$ has been obtained, in contrast with standard ABC approaches. The sensitivity analysis in Table 1 indicates that the vague prior ($a = 100$) does not affect the conclusions, with differences in the posterior only emerging for very strong prior beliefs ($a = 0.1$).

For both spherical and ellipsoidal inclusion models we additionally constructed a likelihood estimate by simply fitting a mixture of multivariate normals to the samples $(\mathbf{S}^{(i)}, \boldsymbol{\theta}^{(i)})$ and then conditioning on $\boldsymbol{\theta}$. This is equivalent to our regression density approach without the transformation $(\mathbf{S}, \boldsymbol{\theta}) \rightarrow (\mathbf{U}, \boldsymbol{\theta})$, or the method of Bonassi et al. (2011) but conditioning on $\boldsymbol{\theta}$ to obtain a likelihood, rather than on \mathbf{S} to produce a posterior. For both models this approach performed reasonably well when the 3 semi-automatic statistics were used (results not shown). However, it became numerically unstable in the 115 dimensional setting. That the regression density approach is numerically stable even in high dimensions indicates the importance of the flexible mixture of experts transformation, $(\mathbf{S}, \boldsymbol{\theta}) \rightarrow (\mathbf{U}, \boldsymbol{\theta})$, in this setting.

5 Discussion

Approximate Bayesian computation using an estimated likelihood function has many advantages over direct estimation of the posterior. As such, we argue that methods that focus on likelihood construction should be considered an important part of the ABC toolbox. While conditional density estimation in high dimensions is challenging, it is an active area of current research. As the reliability and flexibility of these techniques improve, this will only enhance the value of the likelihood-focused approach to ABC.

References

- Beaumont, M. A. (2010). Approximate Bayesian computation in evolution and ecology. *Annu. Rev. Ecol. Evol. Syst.* 41, 379–406.
- Beaumont, M. A., W. Zhang, and D. J. Balding (2002). Approximate Bayesian computation in population genetics. *Genetics* 162, 2025 – 2035.
- Blum, M. G. B. and O. Francois (2010). Non-linear regression models for approximate Bayesian computation. *Statistics and Computing* 1(20), 63 – 73.
- Blum, M. G. B., M. A. Nunes, D. Prangle, and S. A. Sisson (2012). A comparative review of dimension reduction methods in approximate Bayesian computation. *Statistical Science*, in press.
- Bonassi, F. V., L. You, and M. West (2011). Bayesian learning from marginal data in bionetwork models. *Statistical Applications in Genetics and Molecular Biology* 10(1).
- Bortot, P., S. G. Coles, and S. A. Sisson (2007). Inference for stereological extremes. *Journal of the American Statistical Association* 102, 84–92.
- Csilléry, K., M. G. B. Blum, F. Gaggiotti, and O. François (2010). Approximate Bayesian computation (ABC) in practice. *Trends in Ecology and Evolution* (25), 410–418.
- Diggle, P. J. and R. J. Gratton (1984). Monte Carlo methods of inference for implicit statistical models. *Journal of the Royal Statistical Society, Series B* 46(2), 193 – 227.
- Fearnhead, P. and D. Prangle (2012). Constructing summary statistics for approximate Bayesian computation: Semi-automatic approximate Bayesian computation (with discussion). *Journal of the Royal Statistical Society, Series B* 74(3), 419–474.
- Giordani, P., X. Mun, M.-N. Tran, and R. Kohn (2012). Flexible multivariate density estimation with marginal adaptation. *Journal of Computational and Graphical Statistics*, in press.
- Jacobs, R., M. Jordan, S. Nowlan, and G. Hinton (1991). Adaptive mixtures of local experts. *Neural Computation* 3, 79 – 87.
- Jordan, M. and R. Jacobs (1994). Adaptive mixtures of local experts. *Neural Computation* 6, 181 – 214.
- Leuenberger, C. and D. Wegmann (2010). Bayesian computation and model selection without likelihoods. *Genetics* 184, 243 – 252.
- Luciani, F., S. A. Sisson, H. Jiang, A. Francis, and M. M. Tanaka (2009). The high fitness cost of drug resistance in mycobacterium tuberculosis. *Proc. Natl. Acad. Sci. USA* 106, 14711–14715.
- Norets, A. and J. Pelenis (2012). Bayesian modeling of joint and conditional distributions. *Journal of Econometrics* 168, 332–346.

- Nott, D. J., Y. Fan, L. Marshall, and S. A. Sisson (2012). Approximate Bayesian computation and Bayes linear analysis: towards high-dimensional ABC. *Journal of Computational and Graphical Statistics*, in press.
- Nott, D. J., S. L. Tan, M. Villani, and R. Kohn (2012). Regression density estimation with variational methods and stochastic approximation. *Journal of Computational and Graphical Statistics* 21, 797–820.
- Pritchard, J. K., M. T. Seielstad, A. Perez-Lezaun, and M. W. Feldman (1999). Population growth of human Y chromosomes: A study of Y chromosome microsatellites. *Molecular Biology and Evolution* 16, 1791–1798.
- Ratmann, O., C. Andrieu, T. Hinkley, C. Wiuf, and S. Richardson (2009). Model criticism based on likelihood-free inference, with an application to protein network evolution. *Proc. Natl. Acad. Sci. USA* 106, 10576–10581.
- Ratmann, O., O. Jorgensen, T. Hinkley, M. Stumpf, S. Richardson, and C. Wiuf (2007). Using likelihood-free inference to compare evolutionary dynamics of the protein networks of *H. Pylori* and *P. falciparum*. *PLoS Comp. Biol.* 3, e230.
- Robert, C. P., J.-M. Cornuet, J.-M. Marin, and N. Pillai (2011). Lack of confidence in approximate Bayesian computation model choice. *Proc. Natl. Acad. Sci.* 108(37), 15112–15117.
- Sisson, S. A. and Y. Fan (2011). Likelihood-free Markov chain Monte Carlo. In S. P. Brooks, A. Gelman, G. Jones, and X.-L. Meng (Eds.), *Handbook of Markov Chain Monte Carlo*, pp. 319–341. Chapman and Hall/CRC Press.
- Sisson, S. A., Y. Fan, and M. M. Tanaka (2007). Sequential Monte Carlo without likelihoods. *Proc. Natl. Acad. Sci.* 104, 1760–1765. Errata (2009), 106:16889.
- Sisson, S. A., G. W. Peters, Y. Fan, and M. Briers (2008). Likelihood-free samplers. Technical report, University of New South Wales.
- Tavaré, S., D. J. Balding, R. C. Griffiths, and P. Donnelly (1997). Inferring coalescence times from DNA sequence data. *Genetics* 145(505-518).
- Tran, M.-N., D. J. Nott, and R. Kohn (2012). Simultaneous variable selection and component selection for regression density estimation with mixtures of heteroscedastic experts. *Electronic Journal of Statistics* 6, 1170–1199.
- Wasserman, L. A. (1989). A robust Bayesian interpretation of likelihood regions. *The Annals of Statistics* 17(3), 1387 – 1393.
- Wood, S. N. (2010). Statistical inference for noisy nonlinear ecological dynamic systems. *Nature* 466, 1102 – 1104.




Figure 1: Pairwise scatterplots between the three semi-automatic summary statistics S^1, S^2, S^3 (Fearnhead and Prangle, 2012) and the three parameters λ, σ, ξ for the ellipsoidal inclusions model.




Figure 2: Pairwise scatterplots between the mixture of experts transformed statistics, U^1, U^2, U^3 and the three parameters λ, σ, ξ for the ellipsoidal inclusions model.

	Spherical			Ellipsoidal		
	0.025	mean	0.975	0.025	mean	0.975
λ_{mle}	24.25	31.67	39.08	69.53	96.77	124.01
λ_{ABC}	24.26	30.61	37.93	65.01	95.18	135.11
λ_{100}	25.00	31.59	39.29	68.82	93.40	119.38
λ_1	25.09	31.65	39.34	69.41	93.66	118.28
$\lambda_{0.5}$	25.05	31.63	39.32	70.42	93.97	116.76
$\lambda_{0.1}$	25.18	31.77	39.44	72.04	94.98	117.84
σ_{mle}	1.01	1.48	1.94	1.08	1.84	2.60
σ_{ABC}	1.02	1.45	1.95	1.12	1.92	3.06
σ_{100}	1.03	1.46	1.95	1.08	1.88	2.96
σ_1	1.03	1.46	1.96	1.07	1.87	2.95
$\sigma_{0.5}$	1.03	1.46	1.95	1.09	1.85	2.86
$\sigma_{0.1}$	1.07	1.44	1.86	1.11	1.68	2.34
ξ_{mle}	-0.21	-0.03	0.15	-0.29	-0.08	0.14
ξ_{ABC}	-0.18	-0.02	0.17	-0.35	-0.09	0.14
ξ_{100}	-0.18	-0.02	0.15	-0.33	-0.09	0.12
ξ_1	-0.18	-0.02	0.15	-0.33	-0.08	0.12
$\xi_{0.5}$	-0.17	-0.02	0.15	-0.31	-0.08	0.12
$\xi_{0.1}$	-0.13	-0.01	0.12	-0.18	-0.04	0.10

Table 1: Posterior summaries for $\theta = (\lambda, \sigma, \xi)$ for the spherical (left) and ellipsoidal (right) models. Columns correspond to the estimated 0.025 and 0.975 quantiles and the posterior mean for each parameter, under each model. The *mle* parameter suffix indicates results from maximum likelihood estimation (giving the MLE and lower and upper 95% confidence limits), *ABC* indicates estimates from the ABC posterior using the approach of Bortot et al. (2007) with $\delta = 3.3$. A numerical suffix (e.g. $\lambda_a, \sigma_a, \xi_a$) indicates the prior for ξ is $N(0, a^2)$ under the regression density approach.




Figure 3: Mixture of experts model fits of $f_j(S^j|\boldsymbol{\theta})$ for $j = 5, 40, 80$ and 110 . Solid line indicates fitted density at the point $\boldsymbol{\theta}_0 = (\lambda_0, \sigma_0, \xi_0) = (30, 1.5, -0.05)$. Histogram denotes the true density obtained by drawing samples from $L(S^j|\boldsymbol{\theta}_0)$ at the point $\boldsymbol{\theta}_0$.




Figure 4: Estimated marginal posteriors for the spherical inclusions model. Solid lines correspond to true marginal density, obtained using standard MCMC. Histograms indicate marginal posteriors based on regression density estimation using: (top row) 112 summary statistics and (middle row) the 3 semi-automatic summary statistics (Fernhead and Prangle, 2012). Bottom row shows quantile-quantile plots of the 3 summary statistic approximation to the true marginal distributions.

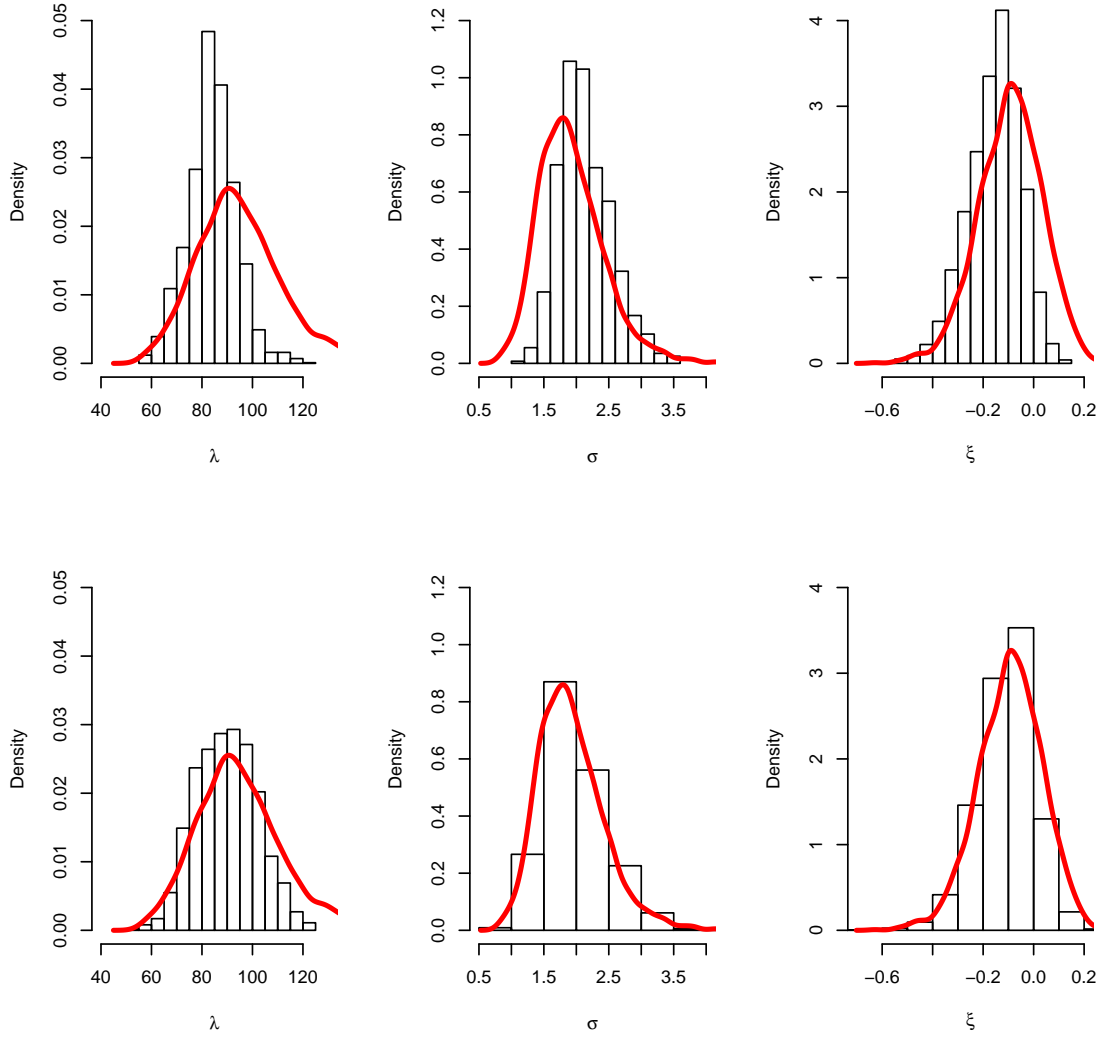


Figure 5: Estimated marginal posteriors for the ellipsoidal inclusions model. Solid lines correspond to the marginal density obtained by the ABC-MCMC approach of Bortot et al. (2007), with kernel scale parameter $\delta = 0.33$. Histograms indicate marginal posteriors based on regression density estimation using: (top row) 112 summary statistics and (bottom row) the 3 semi-automatic summary statistics.